

Economic Viability of Paris Metro Pricing for Digital Services*

Chi-Kin Chau¹, Qian Wang², and Dah-Ming Chiu³

¹Masdar Institute of Science and Technology

²Alibaba, Inc.

³Chinese University of Hong Kong

July 9, 2015

Abstract

Nowadays digital services, such as cloud computing and network access services, allow dynamic resource allocation and virtual resource isolation. This trend can create a new paradigm of flexible pricing schemes. A simple pricing scheme is to allocate multiple isolated service classes with differentiated prices, namely Paris Metro Pricing (PMP). The benefits of PMP are its simplicity and applicability to a wide variety of general digital services, without considering specific performance guarantees for different service classes. The central issue of our study is whether PMP is economically viable, namely whether it will produce more profit for the service provider and whether it will achieve more social welfare. Prior studies had only considered specific models and arrived at conflicting conclusions. In this article, we identify unifying principles in a general setting and derive general sufficient conditions that can guarantee the viability of PMP. We further apply the results to analyze various examples of digital services.

Keywords: Internet Economics; Pricing; Service Classes; Cloud Computing Services

1 Introduction

The management of digital services, such as cloud computing, video streaming and gaming services, and network access services in wireline and wireless networks, is increasingly dictated by economic principles. Digital services possess several salient characteristics that allow more flexible resource management and allocation mechanisms that pave the way for a variety of innovative pricing schemes.

First, there emerge new technologies for virtual resource isolation, by which resources can be allocated conveniently (either dynamically or a priori) to create isolated service classes without altering the underlying hardware infrastructure. Virtual resource isolation can be realized by a variety of technologies. For example, computational virtualization can create virtual machines (VMs) to abstract the underlying hardware resources from applications. Each VM can execute isolated computing tasks without affecting others. On the other hand, software-defined radio has been utilized at base stations to enable dynamic spectrum allocation for different standards in wireless access networks, whereas multiprotocol label switching (MPLS) can assign different routing and queuing operations for packets depending on specific applications. With virtual resource isolation, multiple service classes can be conveniently established.

*This paper appears in ACM Transactions on Internet Technology (ToIT), Special Issue on Pricing and Incentives in Networks and Systems, Vol. 14, No. 12, Issue 2-3, pp12:1-12:21, Oct 2014. A preliminary version has been presented at IEEE INFOCOM 2010 [2].

Second, the availability of real-time measurement of services enables users to observe the performance of other service classes. Users may make instantaneous decisions and adjustment to their behavior according to the observed performance. In the presence of multiple service classes, users who cannot tolerate the performance in the respective class may choose to switch between classes as a consequence. Third, pricing models are evolving so as to be more flexible. The rise of on-demand pricing models, such as pay-per-usage, allows users to readily opt out of a service.

Therefore, it is natural to consider pricing schemes for multiple service classes in digital services. Nowadays, differentiated pricing with multiple service classes has been observed in practice. For example, streaming and file sharing service providers offer premier and economic service classes that can run simultaneously on the same cloud computing platform with isolated VMs and partitioned bandwidth. In mobile wireless service, 3G and 4G can potentially share a spectrum allocated dynamically at base stations. A simple pricing scheme is to impose differentiated prices at different service classes without fulfilling specific quality-of-service requirements and let users spontaneously opt for the appropriate service classes according to their experienced performance. For example, a premier service class charged at a higher price will attract fewer low-end users and therefore is able to provide superior performance to the high-end users. This gives rise to so-called Paris Metro Pricing (PMP), a simple multiclass flat-rate pricing scheme.

PMP is an attractive approach due to its simplicity. However, a more subtle question is whether PMP is economically viable and specifically, *whether we can improve the profit and social welfare through a suitable pricing scheme on service classes with an appropriately allocated amount of resources*. Prior studies arrived at conflicting conclusions for this question. On the one hand, [4, 6] found PMP to be viable, while on the other, based on a similar model, [10] numerically showed that PMP may not be more viable than flat-rate pricing.

In this article, we observe unifying principles that depend on the nature of the externality of the underlying service models. [4, 6] assumed one type of congestion function, whereas [10] assumed another, hence reaching conflicting conclusions. We consider a general model that can capture a wide variety of digital services and we provide general sufficient conditions for the viability of PMP in terms of both social welfare as well as provider profit. This leads to the insights on *why* PMP is or is not viable, understandable by common practitioners. We further apply our results to analyze various examples of digital services.

Outline: Sec. 2 provides the background and related work. Sec. 3 formulates a general model of PMP, and gives several examples of digital services that it captures. Sec. 4 gives an overview of our analytic results considering monopoly. We then extend our study to duopoly, supported by numerical studies in Sec. 5.

2 Background and Related Work

PMP was first proposed by Odlyzko [8] as a simple pricing model for the Internet-differentiated services and can better satisfy users with different aversion to effects of congestion as divided by different service classes. The scheme is inspired by the convention used by Paris metro at one time¹: The first- and second-class cars are charged with different prices, although physically the cars are the same (in terms of the number and quality of the seats). Since fewer people would pay more for the first-class fare, it is also less congested. Thus, users more concerned about getting a seat can opt for first class, and more cost conscious users can opt for second. Note that, in each class, the user is still paying a flat rate. PMP has a self-stabilizing property, namely that if the performance of

¹This scheme is actually adopted rather widely in other transportation systems in the world, including the Mass -Transit Railway in Hong Kong.

first class deteriorates, some users will switch to second class, thus increasing the quality differential between the classes.

In general, the main types of pricing schemes adopted in practice include²: (1) *flat-rate pricing*, that is, to charge a one-off payment for every user, regardless her usage. (2) *usage-based pricing*, that is, to charge according to the amount and pattern of usage of each individual user. In this broad classification, we also include congestion pricing as a form of usage-based pricing. In the presence of multiple service classes, one can apply flat-rate or usage-based pricing to each of the service classes. Congestion pricing, levied only when resource demand exceeds supply, can be argued as economically and theoretically the most optimal strategy for allocating congested resources. In practice, however users strongly prefer flat-rate pricing [7]³.

Since PMP relies on the spontaneous economics adjustment of users, the outcome may not align with the goals of mechanism designers. There have been studies [4–6, 10, 11, 13] in the literature to address the viability of PMP that consider the following two major criteria: (1) *Social welfare*, the total utility of users, taking into account the congestion. (2) *Provider profit*, the total payment collected by the service providers from different service classes.

Based on a multiproduct economics model in [1, 3], [4] formulates a model of PMP, considering the massive number of infinitesimal users and analyzes the viability of PMP for social welfare and provider profit. However, [4] consider only a specific model (that we call utilization-sensitive service) and assume no user will opt out⁴. Next, [5] consider a less specified model but focusing on provider profit and still under the assumption of no user opt-out and homogeneous service classes (where the resources allocated to each class are identical).

On the other hand, [6, 11] consider utilization-sensitive service with user opt-out, but only study provider profit. Meanwhile, [10] numerically studies PMP under a different model (which we call latency-sensitive service) for only provider profit. In the prior work of utilization-sensitive service ([4, 6, 11]), it is reported that PMP is viable for a single monopoly provider. Nonetheless, [10] reports that PMP is not viable in latency-sensitive service for a monopoly provider based on only numerical analysis. We observe that different models can give contradictory results. But there are yet any studies to provide a complete picture considering a general model, beyond the specific examples of utilization- and latency-sensitive services.

To provide a unifying picture, we first analytically study the case of single monopoly provider. We then extend to a duopoly setting for different models. In contrast to the observation in [4], we find that PMP is also viable for duopoly for a certain model of congestion with user opt-out. Our analysis is supported by extensive numerical study. Our study follows the popular model of *infinitesimal users*, as in the prior work ([4–6, 10, 11]). We note that a different model of PMP with finitely many users has been studied in [13] showing that in a specific setting of user utility function wherein a single service class is strictly better than multiple service classes for a monopoly provider. Their conclusion agrees with [10] and generally with ours, albeit by a different model with finitely many users.

²One can also apply a static (time-invariant) pricing strategy as well as a dynamic pricing strategy (that is contingent on the history) to these pricing schemes. Also, the pricing strategy may be application-specific [12].

³ For example, while the Internet service provider (ISP) settlements on aggregate transited traffic are typically based on usage-based pricing, ISPs charge users using flat-rate pricing for its simplicity. This observation also applies to other digital services. While flat-rate pricing is simple, it is insufficient to control the desirable performance to satisfy users' utility. Users are intolerant of inferior performance are forced to opt out. On the other hand, usage-based pricing enables more sophisticated control of performance, but incurs a higher implementation cost because of its more intrusive monitoring and policing on the usage pattern.

⁴Allowing user opt-out leads to a variable number of users in the system that is equivalent to elastic demand, whereas inelastic demand is equivalent to a constant fixed number of users in the system.

3 Model and Notations

This section presents a model for PMP based on different possible forms of negative externality, generalizing the models from [4–6, 10, 11], and defines the notations of *equilibrium* (in which users settle their selections of service classes) and *social welfare* and *provider profit*.

Regardless of the technologies of resource isolation in digital services, it is vital to understand the economical viability of such a service model. Without specifying the implementation details of resource isolation, we assume that resources can be conveniently split or merged among different service classes.

3.1 Utility and Services Classes

Suppose that there are m service classes. Similar to many economic studies [4, 6, 11], we assume there are a large number of users. This can be approximated by a continuum model of *infinitesimal users* such that the type of each infinitesimal user is characterized by a one-dimensional valuation of a positive real parameter θ . The consideration of infinitesimal user is widely used in economics literature that concerns massive numbers of users, such as in digital services. One expects that the more detailed model of finitely many users will approach the simpler one of infinitesimal users when the number of users becomes large.

When a user of type θ (in short, we call user θ) uses service class $i \in \{1, \dots, m\}$, we assume that its utility is given by:

$$U_\theta(i) \triangleq V - p_i - \theta \cdot K(Q_i, C_i) \quad (1)$$

These parameters are explained as follows.

1. V is the maximum utility of accessing the service.
2. p_i is the one-off payment charged per user when accessing the i -th service class. Without loss of generality, we assume $V \geq p_1 \geq p_2 \geq \dots \geq p_m \geq 0$.
3. C_i is the proportion of total capacity of the i -th service class, such that $\sum_{i=1}^m C_i = 1$.
4. Q_i is the volume of users of accessing the i -th service class.
5. K is a congestion function that is increasing in Q_i , but is decreasing in C_i , to be explained in the next section.

Therefore, user θ will have two options either: (1) to select the i -th service class to join that gives the highest utility as

$$i = \arg \max_{j \in \{1, \dots, m\}} U_\theta(j), \quad (2)$$

or (2) to opt out of all service classes because joining any service class will result in a negative utility, that is, $U_\theta(i) < 0$ for all $i \in \{1, \dots, m\}$.

An immediate implication of utility function $U_\theta(i)$ (Eqn. (1)) is that, the larger the value of θ means the higher the valuation on the negative externality, as compared with the price of each service class. Hence, users with the larger value of θ will be more likely to opt out of the service because of negative utility $U_\theta(i) < 0$ for all i .

Remark: Utility function can be alternately defined as $\tilde{U}_\theta(i) \triangleq \theta \cdot \tilde{K}(Q_i, C_i) - p_i$, where $\tilde{K}(Q_i, C_i)$ represents a satisfaction function capturing the inverse effect of a congestion function⁵. We consider utility function as Eqn. (1) in this article, because there are special cases used in the prior work in

⁵In this case, users with smaller value of θ will be likely to opt out of the service because of negative utility. Previously, $\tilde{U}_\theta(i)$ was used in the literature of economics [1, 3].

Table 1: Examples of Metrics of Negative Externality for Digital Services

Utilization	Latency	Loss Probability	Outage Probability
$K_{\text{util}}(Q, C) \triangleq \frac{Q}{C}$	$K_{\text{lat}}(Q, C) \triangleq \frac{1}{C-Q}$	$K_{\text{los}}(Q, C) \triangleq \left(\frac{Q}{C}\right)^k \frac{1-\frac{Q}{C}}{1-(\frac{Q}{C})^{k+1}}$	$K_{\text{out}}(Q, C) \triangleq \left(\frac{Q}{C}\right)^C$

networking research community [4-6, 10, 11]. Nonetheless, the results derived in this article can be easily adapted and similarly applied to $\tilde{U}_\theta(i)$. Also, we remark that [1, 3] consider only homogeneous service classes (i.e., C_i is the same in all m service classes). We relax this constraint to include heterogeneous service classes of different C_i 's.

3.2 Congestion Functions of Digital Service Models

Under PMP, users first observe the performance of different service classes and then opt for the appropriate classes. The performance of digital services is characterized by *negative externality*, where, the greater number of users accessing a certain service, the less favorable performance the users can perceive. The degree of negative externality also depends on the resource allocated to the service. Hence, negative externality can be characterized by: (1) the volume of users accessing the service, and (2) the amount of allocated resource. These two quantities form the basis of negative externality.

In this article, we employ a congestion function $K(Q, C)$ to capture the negative externality, where C represents the numerical amount of resource and where Q represents the volume of users accessing the service. We normalize Q , such that $0 \leq Q \leq C$. Table 1 provides several examples of $K(Q, C)$ that can model various metrics of negative externality in common digital services.

1. *Utilization*: This is measured by the portion of capacity per each unit of load, given by $K_{\text{util}}(Q, C) \triangleq \frac{Q}{C}$. Utilization is a useful metric for computation, bandwidth, and memory sharing services, such as cloud computing and network access services. $K_{\text{util}}(Q, C)$ was considered in prior work [4].
2. *Latency*: This is conveniently captured by a simple M/M/1 queue. Assuming the arrival rate is of Q units, and the service rate is of C units, then the total expected waiting time (i.e., queuing time plus service time) is $K_{\text{lat}}(Q, C) \triangleq \frac{1}{C-Q}$. $K_{\text{lat}}(Q, C)$ was considered in prior work [10]. Latency is a useful metric for queuing sensitive services such as streaming or network routing services. We also consider more general M/G/1 queue in the later section.
3. *Loss Probability*: We also consider M/M/1/ k queue to model queuing-based digital services. An important metric is whether a request will be dropped when all k servers are occupied. In the M/M/1/ k model, the probability that k servers are occupied is given by $K_{\text{los}}(Q, C) \triangleq \left(\frac{Q}{C}\right)^k \frac{1-\frac{Q}{C}}{1-(\frac{Q}{C})^{k+1}}$.
4. *Outage Probability*: To model the reliability of digital services, we consider a small probability that a server will fail. A natural setting of failure probability is proportional to the utilization $\frac{Q}{C}$. If there are C servers, then the probability that all servers will fail is given by $K_{\text{out}}(Q, C) \triangleq \left(\frac{Q}{C}\right)^C$.

Observations: Different congestion functions lead to different results. As a clear illustration, we consider a simple scenario of resource partitioning into two identical service classes charged at an identical price such that each class is allocated with half the original resource. Because of identical resource and price, it is likely the usage will be split equally between the two identical service classes. Hence, we define the resource and usage for each service class by $C_1 = C_2 = \frac{C}{2}$ and $Q_1 = Q_2 = \frac{Q}{2}$, where the subscript indicates the first or second class, and C, Q are the resource and usage of the original nonsplit service class. For utilization-sensitive service, the congestion functions after service partitioning become

$$K_{\text{util}}(Q_1, C_1) = K_{\text{util}}(Q_2, C_2) = \frac{Q}{C} = K_{\text{util}}(Q, C) \quad (3)$$

In other words, the users will not perceive any difference in terms of negative externality. On the other hand, for latency-sensitive service, the congestion functions after service partitioning become

$$K_{\text{lat}}(Q_1, C_1) = K_{\text{lat}}(Q_2, C_2) = \frac{2}{C - Q} > K_{\text{lat}}(Q, C) \quad (4)$$

In this case, the users actually perceive a degradation of service after service partitioning! Therefore, we may conclude that the provision of multiple service classes provided by resource partitioning is not viable for latency-sensitive service, because of the decrease of total welfare of users and also the decrease of provider profit, as some users may opt out of the service. The preceding simple discussion, however, has not taken into consideration that different service classes may have varying prices. A more complete study requires a model including payment as a part of the user utility and a clear description of the users' decision process for switching between service classes and opting out. This will be completed in the following sections.

3.3 Equilibrium, Social Welfare and Provider Profit

In this article, we will analyze the viability of PMP at equilibrium. An *equilibrium* will be attained when no user switches from his selection. To formulate equilibrium, we first note that $K(Q_i, C_i)$ is fixed for each i -th service class at equilibrium, thus also the price p_i . Hence, the utility function $U_\theta(i)$ (Eqn. (1)) becomes a linear function of θ with $-K(Q_i, C_i)$ as the slope and $V - p_i$ as the y -intercept. For the case $m = 2$, we plot the utility $U_\theta(i)$ against θ in Fig. 1 for illustration.

Given a vector of differentiated prices, denoted by $\mathbf{p} = (p_i)_{i=1}^m$, for the m service classes, there exists a set of *cut-off users*, denoted by $\boldsymbol{\theta} = (\theta_i)_{i=1}^m$, such that, for $i = 2, \dots, m$, each cut-off user θ_i is indifferent to joining the $(i-1)$ -th service class or the i -th service class (i.e., $U_{\theta_i}(i-1) = U_{\theta_i}(i)$), and cut-off user θ_1 is indifferent to joining the first service class or opting out of these m service classes. Hence, an equilibrium can be characterized by the tuple $(\mathbf{p}, \boldsymbol{\theta})$.

Let $Q_0 \triangleq 1 - \sum_{i=1}^m (Q_i)$ be the volume of users who opt out of these m service classes. The users are parameterized by θ , which is described by a cumulative distribution function $F(\theta)$ (and its probability density function denoted by $f(\theta)$). A special case is that $F(\theta)$ is a uniform distribution.

Definition 3.1. (Equilibrium) *The tuple $(\mathbf{p}, \boldsymbol{\theta})$ defines an equilibrium, if the following constraints are satisfied:*

$$(c.1): \theta_1 > \theta_2 > \dots > \theta_m > \theta_{m+1} = 0,$$

$$(c.2): K(Q_i, C_i) \leq K(Q_{i+1}, C_{i+1}) \text{ for } i \neq m, \text{ where}$$

$$Q_i \triangleq \begin{cases} F(\theta_m) & \text{if } i = m \\ F(\theta_i) - F(\theta_{i+1}) & \text{if } 1 \leq i < m \end{cases} \quad (5)$$




Figure 1: An illustration of equilibrium. We plot the utility $U_\theta(i)$ against θ for the case $m = 2$. We assume $F(\theta) = \theta$ is a uniform distribution.

(c.3):

$$\begin{aligned} p_{i-1} - p_i &= \theta_i \cdot (K(Q_i, C_i) - K(Q_{i-1}, C_{i-1})) & \text{if } 1 < i \leq m \\ p_1 &= V - \theta_1 \cdot K(Q_1, C_1) & \text{otherwise} \end{aligned} \quad (6)$$

In other words, (c.1) requires the set of cut-off users $(\theta_i)_{i=1}^{m+1}$ to have a strict order of valuations on negative externality. User $\theta_{m+1} = 0$ is the least-valuation user who always accepts the lowest-priced service class (i.e., the m -th service class). (c.2) follows from $V \geq p_1 \geq \dots \geq p_m \geq 0$ and the definition of utility function (Eqn. (1)). Lastly, (c.3) characterizes the prices of service classes based on the fact that each cut-off user θ_i is indifferent to joining the $(i-1)$ and i -th service classes (i.e., $U_{\theta_i}(i-1) = U_{\theta_i}(i)$). These conditions are depicted in Fig. 1.

Note that it is possible that $p_i = p_{i-1}$ (i.e., $K(Q_i, C_i) = K(Q_{i-1}, C_{i-1})$) for some i . It is easy to see that an equilibrium always exists given either \mathbf{p} or $\boldsymbol{\theta}$. Namely, there is a one-to-one mapping between \mathbf{p} and $\boldsymbol{\theta}$ at equilibrium (see [1] for a rigorous proof).

To facilitate the analysis, we also assume that $F(\theta)$ is a well-formed distribution, such that $f(\theta) > 0$ for $\theta \in [0, \bar{\theta}] \subseteq [0, 1]$ for constant $\bar{\theta}$, and $f(\theta) = 0$ otherwise. This assumption prevents discontinuity of the marginal change of \mathbf{p} with respect to the marginal change of $\boldsymbol{\theta}$ at equilibrium.

Finally, we define the social welfare as the total utility of all users excluding the payment and define the monopoly provider profit as the total payment collected from the users. Suppose $(\mathbf{p}, \boldsymbol{\theta})$ is an equilibrium. Let the *social welfare* be

$$S(\mathbf{p}) \triangleq \sum_{i=1}^m \int_{\theta_{i+1}}^{\theta_i} (V - \theta \cdot K(Q_i, C_i)) \cdot f(\theta) d\theta \quad (7)$$

Note that in Eqn. (7), $\boldsymbol{\theta}$ is a function of \mathbf{p} at equilibrium. Let the monopoly *provider profit* be

$$\pi(\mathbf{p}) \triangleq \sum_{i=1}^m p_i \cdot Q_i \quad (8)$$

3.4 Numerical Study

Before presenting our analytical results, this section provides some illustrations of the consequences of PMP in terms of social welfare and provider profit at equilibrium by means of numerical studies.

We consider four specific examples, namely utilization-, latency-, loss- and outage-sensitive services. Specifically, we compare: (1) the case of one single service class ($C = 1$), and (2) the case of two service classes ($C_1 = 0.3, C_2 = 0.7$). In the numerical example, we let the probability distribution $F(\theta)$ be a uniform distribution and the maximum utility be $V = 2$.

We compare the maximum social welfare and provider profit that can be achieved between one single service class and two service classes. For two service classes, we let $p_2 = a \cdot p_1$, where $a \in [0, 1]$. When $a = 1$, it is equivalent to identical pricing at two service classes. We numerically evaluated the following four quantities for different service models.

$$\max_{0 \leq p \leq V} \pi(p), \quad \max_{0 \leq p_1 \leq V} \pi(p_1, a \cdot p_1), \quad \max_{0 \leq p \leq V} S(p), \quad \max_{0 \leq p_1 \leq V} S(p_1, a \cdot p_1) \quad (9)$$

Figs. 2-5 show the respective numerical values of maximum social welfare and provider profit at given values of a , for utilization-, latency-, loss- and outage-sensitive services. Note that, when there is one service class, the maximum social welfare and provider profit are constant values with respect to a .

Observations: We obtain the following observations from Figs. 2-9 that motivate our later results.

1. For utilization- and loss-sensitive services, the social welfare and provider profit for one service class and two service classes under identical pricing (i.e., $a = 1$) are equivalent. But for latency- and outage-sensitive service, there is a gap between the cases, that is, one service class yields better social welfare and profit for latency-sensitive service, whereas it is worse for outage-sensitive service.
2. For utilization-, loss- and outage-sensitive services, the service provider can make more profit from PMP under a certain setting of differentiated pricing (see Figs 3, 7, 9); but for latency-sensitive service, PMP (whether under differentiated pricing or identical pricing) is not viable from the provider profit's point of view (see Fig 5). These are exactly the contradictory conclusions reached by [6] and [10], respectively. We observe the same conclusions for social welfare.
3. We remark that, in general, for latency-sensitive service with two service classes, identical pricing may or may not provide higher social welfare or profit than differentiated pricing. However, for utilization-, loss- and outage-sensitive services with two service classes, it is always possible that differentiated pricing can provide higher social welfare and profit than identical pricing.

The numerical results just given motivate the goal for this study, namely how to settle the question of when PMP can be guaranteed to yield higher profits and achieve more social welfare. We answer this question by deriving conditions on a general class of congestion functions, beyond just the cases of utilization- and latency-sensitive services.

Tactically, we study this problem in two steps by asking the following questions.

1. How do we ensure the viability of partitioning a service class into multiple service classes under identical pricing?
2. How do we ensure the viability of differentiated pricing of multiple service classes as compared with identical pricing?

Resolving these two questions can lead to a characterization of the viability of PMP. Our answers to these questions are elaborated in Sec. 4.




Figure 2: Maximum social welfare for utilization-sensitive service, where $p_2 = a \cdot p_1$.




Figure 3: Maximum provider profit for utilization-sensitive service, where $p_2 = a \cdot p_1$.




Figure 4: Maximum social welfare for latency-sensitive service, where $p_2 = a \cdot p_1$.



Figure 5: Maximum provider profit for latency-sensitive service, where $p_2 = a \cdot p_1$.



Figure 6: Maximum social welfare for loss-sensitive service, where $p_2 = a \cdot p_1$.




Figure 7: Maximum provider profit for loss-sensitive service, where $p_2 = a \cdot p_1$.




Figure 8: Maximum social welfare for outage-sensitive service, where $p_2 = a \cdot p_1$.




Figure 9: Maximum provider profit for outage-sensitive service, where $p_2 = a \cdot p_1$.

4 Monopoly Case

The viability of PMP ultimately depends on the basic properties of the congestion function. In this section, we provide a general analytical study by identifying some key properties of congestion functions associated with the viability of PMP considering monopoly.

4.1 Viability of Service Partitioning

We first consider identical pricing and provide insights for the viability of *service partitioning*, wherein one single service class is partitioned into two service classes, both of which are priced the same as the original service class and, each service class is allocated a portion of the capacity of the original service class. Initially, there is one single service class. Let p and C be the price and capacity of the original single service class. Let $\tilde{\theta}$ be the cut-off user at equilibrium (who is indifferent to either opting-out or joining the service) and the total usage be $\tilde{Q} \triangleq F(\tilde{\theta})$. Then, the social welfare becomes

$$S(p) \triangleq \int_0^{\tilde{\theta}} \left(V - \theta \cdot K(\tilde{Q}, C) \right) \cdot f(\theta) d\theta \quad (10)$$

and the respective provider profit becomes

$$\pi(p) \triangleq p\tilde{Q} \quad (11)$$

Next, we consider resource partitioning into two service classes. Let C_1 and C_2 be the respective capacity of each service class, where $C_1 + C_2 = C$. Let θ_1 and θ_2 be the cut-off users at the equilibrium of each service class, and the respective usage be $Q_1 \triangleq F(\theta_1) - F(\theta_2)$ and $Q_2 \triangleq F(\theta_2)$. Because of an identical pricing of each service class at p , the social welfare for the partitioned service classes is⁶:

$$S(p, p) \triangleq \int_0^{\theta_2} \left(V - \theta \cdot K(Q_2, C_2) \right) \cdot f(\theta) d\theta + \int_{\theta_2}^{\theta_1} \left(V - \theta \cdot K(Q_1, C_1) \right) \cdot f(\theta) d\theta \quad (12)$$

while the respective provider profit is⁷:

$$\pi(p, p) \triangleq p(Q_1 + Q_2) \quad (13)$$

We next provide a general sufficient condition on $K(Q, C)$ for the viability of service partitioning.

Theorem 4.1. *For all $0 \leq \alpha < 1$ we have the following.*

1. (Partition-preferred congestion function): *If $K(Q, C) \geq K(\alpha Q, \alpha C)$,*

$$S(p, p) \geq S(p) \text{ and } \pi(p, p) \geq \pi(p) \quad (14)$$

2. (Multiplexing-preferred congestion function): *If $K(Q, C) \leq K(\alpha Q, \alpha C)$,*

$$S(p, p) \leq S(p) \text{ and } \pi(p, p) \leq \pi(p) \quad (15)$$

The proofs are all in the Appendix unless otherwise stated.

Theorem 4.1 classifies two types of congestion functions. Intuitively, the congestion function favoring service partitioning is one that sees decreased congestion externality as we scale down the usage and capacity. Alternatively, the congestion function that favors multiplexing (service merging) is one that sees decreased congestion externality as we scale up the usage and capacity. Note that a congestion function $K(Q, C)$ is given by definition increasing in usage Q , but decreasing in capacity C . When we scale up both usage and capacity, one or the other of these factors is more dominating, giving rise to the two classes of congestion functions.

We next apply Theorem 4.1 to several specific examples.

⁶In trying to reduce notations, we are going to slightly abuse our use of them here. The social welfare for the case of a single class service with one price p , and the case of multiple service classes with price \mathbf{p} will both be denoted by a function $S(\cdot)$, with a different number of price parameters as appropriate.

⁷Again, the same abuse of notation is applied to the profit function.

Example 1. 1. Utilization-sensitive service: $K_{\text{util}}(Q, C)$ is indifferent to partitioning or multiplexing.

$$K_{\text{util}}(Q, C) = \frac{Q}{C} = \frac{\alpha Q}{\alpha C} = K_{\text{cs}}(\alpha Q, \alpha C) \quad (16)$$

2. Latency-sensitive service: $K_{\text{lat}}(Q, C)$ prefers multiplexing to partitioning.

$$K_{\text{lat}}(Q, C) = \frac{1}{C - Q} < \frac{1}{\alpha(C - Q)} = K_{\text{lat}}(\alpha Q, \alpha C) \quad (17)$$

Further, motivated by Pollaczek-Khinchine formula for $M/G/1$ queue, we also consider general latency as

$$K_{\text{glat}}(Q, C) = \frac{Q(1 + \delta_S^2)}{2C(C - Q)} + \frac{1}{C}, \quad (18)$$

where δ_S^2 is the coefficient of variation of service time, by convention. Similarly, $K_{\text{glat}}(Q, C)$ prefers multiplexing to partitioning.

$$K_{\text{glat}}(Q, C) < \frac{1}{\alpha} K_{\text{glat}}(Q, C) = K_{\text{glat}}(\alpha Q, \alpha C) \quad (19)$$

3. Loss-sensitive service: Similar to $K_{\text{util}}(Q, C)$, $K_{\text{los}}(Q, C)$ is indifferent to partitioning or multiplexing.

4. Outage-sensitive service: $K_{\text{out}}(Q, C)$ prefers partitioning to multiplexing:

$$K_{\text{out}}(Q, C) = \left(\frac{\epsilon \alpha Q}{\alpha C}\right)^C > \left(\frac{\epsilon \alpha Q}{\alpha C}\right)^{\alpha C} = K_{\text{out}}(\alpha Q, \alpha C) \quad (20)$$

By means of Theorem 4.1, we conclude that service partitioning is viable for utilization-, loss- and outage-sensitive services. This has been validated in Sec. 3.4. Note that Theorem 4.1 is sufficiently general and can be applied to diverse types of congestion functions.

The following corollary extends Theorem 4.1 from two classes to multiple classes.

Corollary 4.2. Suppose $0 \leq \alpha < 1$, and $\mathbf{p} = (p_i = p)_{i=1}^m$, then the following hold.

1. (Partition-preferred congestion function):

If $K(Q, C) \geq K(\alpha Q, \alpha C)$ for all α ,

$$S(\mathbf{p}) \geq S(p) \text{ and } \pi(\mathbf{p}) \geq \pi(p) \quad (21)$$

2. (Multiplexing-preferred congestion function):

If $K(Q, C) \leq K(\alpha Q, \alpha C)$ for all α ,

$$S(\mathbf{p}) \leq S(p) \text{ and } \pi(\mathbf{p}) \leq \pi(p) \quad (22)$$

Proof. This is straightforward to prove by induction through proving m service classes is true when supposing $(m - 1)$ service classes is true. \square

4.2 Viability of Differentiated Pricing

Although Sec. 4.1 provides the sufficient condition for service partitioning under identical pricing, it does not cover the case of differentiated pricing. In particular, PMP is not perceivable by users under identical pricing. To complete the picture, in this section we compare differentiated pricing and identical pricing. We will rely on the following property of congestion functions.

Definition 4.3. (Monotone Preference to Service Classes) *Given a fixed set $\{C_i : i = 1, \dots, m\}$, the set of congestion functions $\{K(Q_i, C_i) : i = 1, \dots, m\}$ are subject to:*

1. *each $K(Q_i, C_i)$ must be strictly increasing and differentiable in Q_i , hence, the partial derivative of $K(Q_i, C_i)$ at Q_i : $k(Q_i, C_i) \triangleq \frac{\partial K(Q, C)}{\partial Q}|_{Q=Q_i, C=C_i}$ exists and is positive;*
2. *suppose $C_1 < C_2 < \dots < C_m$, then either one of the following two cases must be true:*
 - (m.1) $Q_i > Q_j$ implying $k(Q_i, C_i) > k(Q_j, C_j)$ for any distinct pair $i \neq j$; or
 - (m.2) $Q_i > Q_j$ implying $k(Q_i, C_i) < k(Q_j, C_j)$ for any distinct pair $i \neq j$.

The first condition in Definition 4.3 ensures the smoothness of congestion functions, whereas the second condition requires a monotone order on the derivatives of the m service classes. Intuitively, a monotone order of the derivatives (i.e., (m.1) or (m.2)) reflects a monotone order of sensitivity of negative externality among the service classes. Consequently, such a monotone preference is an indication that, once a user joins the second class, any marginal change to usage in each service class will only shift the user's selection to either the first class or third class (if available). This precludes the switching among service classes that is out of a linear order.

Note that homogeneous service classes with the same convex congestion function (i.e., $C_i = C_j$) obviously satisfy Definition 4.3. Here, we also allow distinct C_i to capture the heterogeneous capacities among different service classes. The provider profit for homogeneous service classes had been studied in [1]. Here we generalize the result to consider heterogeneous service classes and social welfare. We next apply Definition 4.3 to some examples.

Example 2. 1. Utilization-sensitive service: Suppose $C_1 < C_2 < \dots < C_m$, $K_{\text{util}}(Q, C)$ satisfies monotone preference to service classes. Then

$$k_{\text{util}}(Q_i, C_i) = \frac{\partial K_{\text{util}}(Q, C)}{\partial Q}|_{Q=Q_i, C=C_i} = \frac{1}{C_i} \quad (23)$$

$$\Rightarrow k_{\text{util}}(Q_1, C_1) > k_{\text{util}}(Q_2, C_2) > \dots > k_{\text{util}}(Q_m, C_m) \text{ for all } Q_1, Q_2, \dots, Q_m. \quad (24)$$

2. Latency-sensitive service: However, $K_{\text{lat}}(Q, C)$ does not always satisfy monotone preference to service classes because

$$k_{\text{lat}}(Q_i, C_i) = \frac{\partial K_{\text{lat}}(Q, C)}{\partial Q}|_{Q=Q_i, C=C_i} = \frac{1}{(C_i - Q_i)^2} \quad (25)$$

Note that $\frac{1}{(C_i - Q_i)^2}$ is not a monotone function in C_i or Q_i . For instance, when $C_1 = 0.3, C_2 = 0.7, Q_1 = 0.2, Q_2 = 0.5$, then $\frac{1}{(C_1 - Q_1)^2} > \frac{1}{(C_2 - Q_2)^2}$. However, when $Q_1 = 0.05, Q_2 = 0.5$, then $\frac{1}{(C_1 - Q_1)^2} < \frac{1}{(C_2 - Q_2)^2}$.

3. Loss-sensitive service: Similar to $K_{\text{util}}(Q, C)$, $K_{\text{los}}(Q, C)$ satisfies monotone preference to service classes.
4. Outage-sensitive service: One also can show that $K_{\text{out}}(Q, C)$ satisfies monotone preference to service classes.

In the following, we compare the maximum social welfare gained by identical pricing and differentiated pricing, given a monotone preference to service classes. Hence, we conclude that differentiated pricing is viable for utilization-, loss- and outage-sensitive services. This has been validated in Sec. 3.4.

Theorem 4.4. *Given two service classes that satisfy monotone preference (Definition 4.3), the social welfare obtained through identical pricing at p is strictly inferior to differentiated pricing for some $p_1 \neq p_2$:*

$$S(p_1, p_2) > S(p, p) \quad (26)$$

The proof relies on the notion of total derivative $dS(p_1, p_2)$, which can be found in the Appendix (Sec. 7.2). One might not be surprised to see that the differentiated pricing ($p_1 \neq p_2$) could be better than identical pricing ($p_1 = p_2$), but what is remarkable is the strict superiority of differentiated pricing. In Sec. 3.4, we have validated this result for utilization-sensitive service. Note that for a congestion function that does not satisfy monotone preference, the maximum social welfare gained by identical pricing may or may not be higher than that of differentiated pricing. This can be observed in latency-sensitive service. Theorem 4.4 can be extended to the case of m service classes as follows. Hence, for utilization-sensitive service, it is advantageous to offer as many service classes as possible, in terms of an increase of social welfare.

Corollary 4.5. *Given m service classes that satisfy monotone preference (Definition 4.3), let $\mathbf{p} \triangleq (p_i = p)_{i=1}^m$, then there exists $\mathbf{p}' \triangleq (p'_i)_{i=1}^m$, such that $p'_i \neq p'_j$ for all distinct $i, j \in \{1, \dots, m\}$,*

$$S(\mathbf{p}') > S(\mathbf{p}) \quad (27)$$

We can also prove the same results for provider profit.

Theorem 4.6. *Given two service classes that satisfy monotone preference (Definition 4.3), the provider profit of identical pricing at p is strictly inferior to that of differentiated pricing for some $p_1 \neq p_2$:*

$$\pi(p_1, p_2) > \pi(p, p) \quad (28)$$

Corollary 4.7. *Given m service classes that satisfy monotone preference (Definition 4.3), let $\mathbf{p} \triangleq (p_i = p)_{i=1}^m$, then there exists $\mathbf{p}' \triangleq (p'_i)_{i=1}^m$, such that $p'_i \neq p'_j$ for some distinct $i, j \in \{1, \dots, m\}$,*

$$\pi(\mathbf{p}') > \pi(\mathbf{p}) \quad (29)$$

Corollary 4.7 is an immediate consequence of Theorem 4.6, but is weaker than Corollary 4.5, since provider profit is more difficult to analyze than social welfare.

Ramification: We arrive at the point where we can provide a more complete answer to our original question on the viability of PMP. Combining Theorems 4.1-4.6, we have a sufficient condition to guarantee PMP to be viable in the sense of both provider profit as well as social welfare. This also explains not only why PMP is viable for utilization-, loss- and outage-sensitive services, but also for latency-sensitive services is not always viable. Theorems 4.1-4.6 are sufficiently general so that they can also be applied to general digital services.

5 Duopoly Case

The study involving multiple competitive providers is more challenging. A plausible outcome is that the providers settle at a Nash equilibrium at which unilateral change in pricing or capacity allocation

will induce an inferior profit. Previously, [4] studied the properties of Nash equilibrium of PMP in the presence of two competitive providers and derived a closed-form solution of Nash equilibrium for the simple setting of utilization services of the same capacity and *disallowing user opt-out*. They reported that PMP is not viable in the setting of two competitive providers as compared to the case of simply a single service class offered by each provider.

However, in a more general model *with user opt-out* as studied in this article, the viability of PMP is a consequence of more subtle properties of the congestion function. In fact, we find PMP viable in several settings of partition-preferred congestion functions.

In this section, we first derive necessary conditions for Nash equilibrium for two competitive providers. We restrict our analysis to the setting wherein one provider (I) offers a single unpartitioned service class, whereas another provider (II) can flexibly partition its service classes and offer differentiated pricing. The necessary conditions are then applied to the numerical study of several specific congestion functions, from which we compare the viability of PMP.

5.1 Characterization of Nash Equilibrium

We define the profit of provider I as $\pi^I \triangleq p^I Q^I$ where p^I is the price offered for its single service class and Q^I is the amount of users accessing it. From Eqn. (5), $Q^I = Q_i$ if p^I is the i -th highest price among all the service classes offered.

Provider II, by contrast, has an option to offer two service classes (offering $p_1^{II} \geq p_2^{II}$ at capacity C_1^{II} and C_2^{II} respectively), or a single service class (offering p^{II} at $C^{II} = C_1^{II} + C_2^{II}$). Hence

$$\pi^{II} \triangleq \begin{cases} p_1^{II} Q_1^{II} + p_2^{II} Q_2^{II} & \text{for two service classes} \\ p^{II} Q^{II} & \text{for one service class} \end{cases} \quad (30)$$

where $Q_1^{II}, Q_2^{II}, Q^{II}$ are the amount of users accessing the respective service classes.

The necessary condition for Nash equilibrium is that derivative $\frac{d\pi^I}{dp^I} = 0$, at given $p_1^{II}, p_2^{II}, C_1^{II}, C_2^{II}$ (or p^{II}, C^{II}). The respective derivatives are listed for all five cases in Table 2 in the Appendix. In the table, we let $K_i = K(Q_i, C_i)$ and $k_i = \frac{\partial K(Q, C)}{\partial Q} \big|_{Q=Q_i, C=C_i}$, where Q_i is the total amount of users and C_i is the capacity of i -th service class.

Based on the derivatives, we then numerically evaluate p^I such that $\frac{d\pi^I}{dp^I} = 0$ at given $C^I, p_1^{II}, p_2^{II}, C_1^{II}, C_2^{II}$ (or p^{II}, C^{II}). Furthermore, we corroborate the existence of Nash equilibrium at the corresponding $(p^I, C^I, p_1^{II}, p_2^{II}, C_1^{II}, C_2^{II})$ and $(p^I, C^I, p^{II}, C^{II})$ by examining whether π^I and π^{II} are local maxima in the corresponding neighborhood region.

5.2 Numerical Study and Observations

Using the preceding results, we particularly study the viability of PMP for utilization-sensitive service. We generalize the congestion function by incorporating default consumption. Define a modified congestion function as

$$K_{\text{utl.d}}(Q, C) \triangleq \frac{Q - \varepsilon}{C} \quad (31)$$

where $\varepsilon \leq Q$ is a certain consumption incurred whenever the service class is accessed. This is useful to model the scenario wherein certain overhead or default consumption is imposed in the service.

It is straightforward to show that congestion function $K_{\text{utl.d}}(Q, C)$ strictly prefers service partitioning to multiplexing (when $\varepsilon > 0$), and satisfies monotone preferences to service classes.

Figs. 10-13 show the profit of each provider at various given values of p^I , and the corresponding best response from provider II on its pricing and capacity so as to maximize its profit π^{II} .

Observations: We obtain the following key observations from Figs. 10-13:

1. PMP is viable in several settings of utilization service with and without default consumption. Because provider II always has a higher profit when it partitions its service classes (i.e., lines of “2 vs 1” in the figures), so is the profit of provider I.
2. Multiplexing-preferred congestion functions ($K_{\text{utl.d}}(Q, C)$) produce higher viability for PMP.
3. When more capacity is given to provider II, the benefit of PMP is more prominent (see the larger gap between profits of provider I and provider II in Fig. 13).

We note that although general analytical results are harder for the competitive case, our numerical results agree with the observations in other multiplexing-preferred congestion functions.




Figure 10: Utilization-sensitive service ($K_{\text{utl.d}}(Q, C)$, $V = 2$, $C^I = C^{II} = 1$).




Figure 11: Utilization-sensitive service with default consumption ($K_{\text{utl.d}}(Q, C)$, $V = 2$, $C^I = C^{II} = 1$, $\varepsilon = 0.1$).




Figure 12: Utilization-sensitive service with default consumption ($K_{\text{utl.d}}(Q, C)$, $V = 2$, $C^I = C^{II} = 1$, $\varepsilon = 0.2$).




Figure 13: Utilization-sensitive service ($K_{\text{utl}}(Q, C)$, $V = 3$, $C^I, C^{II} = 2$).

6 Conclusion and Discussion

This article provides general conditions for the viability of Paris Metro Pricing (PMP) based on a general setting of negative externality (via a general congestion function) that can model a wide range of digital services. There are two separate messages here, both of which can be intuitively stated. The first one says that one service either prefers multiplexing (having more people share proportionally more capacity) or not; if one wants to guarantee a gain (in terms of profit or social

welfare) by dividing ones service into multiple classes with the same price, then this service would better not prefer multiplexing. The second message says that, if one starts with a multiclass service with the same price and one wants to move to charging different prices, then the service classes one sets up would better generate a monotone linear preference perceived by the users. By combining these two rules together, we can characterize a large class of services that can benefit from PMP. These observations also help us understand why sometimes PMP is not viable. Our results help clarify the confusion caused by conflicting results on the viability of PMP by previous studies. Our model is general and the results are rigorously proved and applicable to future studies on network economics.

References

- [1] Parkash Chander and Luc Leruth. The optimal product mix for a monopolist in the presence of congestion effects. *Intl. Journal of Industrial Organization*, 7(4):437–449, 1989.
- [2] Chi-Kin Chau, Qian Wang, and Dah-Ming Chiu. On the viability of Paris Metro pricing for communication and service networks. In *Proc. IEEE INFOCOM*, 2010.
- [3] A. de Palma and Luc Leruth. Congestion and game in capacity: A duopoly analysis in the presence of network externalities. *Annales d’Economie et de Statistique*, (15–16):389–407, 1989.
- [4] Richard Gibbens, Robin Mason, and Richard Steinberg. Internet service classes under competition. *IEEE Journal Selected Areas in Communications*, 18(2):2490–2498, 2000.
- [5] Ori Haimanko and Richard Steinberg. *Price Symmetry in a Duopoly with Congestion*, December 2000. CORE Discussion Paper.
- [6] R. Jain, T. Mullen, and R. Hausman. Analysis of Paris Metro pricing strategy for QoS with a single service provider. In *Proc. IEEE/IFIP Intl. Workshop on Quality of Service*, pages 495–500, 2001.
- [7] A. Odlyzko, B. St. Arnaud, E. Stallman, and M. Weinberg. now your limits: Considering the role of data caps and usage based billing in internet access service. *Public Knowledge white paper*, 2012.
- [8] Andrew Odlyzko. Paris metro pricing for the internet. In *Proc. ACM Conf. on Electronic Commerce*, pages 140–147, 1999.
- [9] K. F. Riley, M. P. Hobson, and S. J. Bence. *Mathematical Methods for Physics and Engineering*. Cambridge University Press, 2006.
- [10] David Ros and Bruno Tuffin. A mathematical model of the Paris Metro pricing scheme for charging packet networks. *Computer Networks*, 46:73–85, 2004.
- [11] H. Sakurai, S. Kasahara, and N. Adachi. Internet pricing and user opt-out strategy under two ISPs competition. In *Proc. Intl. Network Optimization Conf. (INOC)*, pages 495–500, 2003.
- [12] Soumya Sen, Carlee Joe-Wong, Sangtae Ha, , and Mung Chiang. Incentivizing time-shifting of data: A survey of time-dependent pricing for internet access. *IEEE Communications Magazine November 2012*, 11(50):91–99, 2012.
- [13] S. Shakkottai, R. Srikant, A. Ozdaglar, and D. Accemoglu. The price of simplicity. *IEEE Journal Selected Areas in Communications*, 26(7):1269–1276, 2007.

7 Appendix

7.1 Proof for Theorem 4.1

Proof. First, by Eqn. (6) in Definition 3.1, considering a single service class, we obtain

$$p = V - \tilde{\theta} \cdot K(\tilde{Q}, C) \Rightarrow \tilde{\theta} = \frac{V - p}{K(\tilde{Q}, C)} \quad (32)$$

where $\tilde{\theta}$ is the cut-off user for a single service class. When there are two service classes after resource partitioning, we obtain

$$\begin{cases} p - p = \theta_2 \cdot (K(Q_2, C_2) - K(Q_1, C_1)) \\ p = V - \theta_1 \cdot K(Q_1, C_1) \end{cases} \quad (33)$$

Hence due to Eqn. (33) it follows that

$$K(Q_2, C_2) = K(Q_1, C_1) \quad (34)$$

$$\theta_1 = \frac{V - p}{K(Q_1, C_1)} \quad (35)$$

Using Eqns. (32) and (35), we obtain the following equality.

$$\frac{\tilde{\theta}}{\theta_1} = \frac{K(Q_1, C_1)}{K(\tilde{Q}, C)} = \frac{K(Q_2, C_2)}{K(\tilde{Q}, C)} \quad (36)$$

To complete the proof, we proceed in two steps as follows.

Step 1: Then, we want to show:

1. If $K(Q, C) \geq K(\alpha Q, \alpha C)$ for all $0 \leq \alpha \leq 1$, then $\theta_1 \geq \tilde{\theta}$
2. If $K(Q, C) \leq K(\alpha Q, \alpha C)$ for all $0 \leq \alpha \leq 1$, then $\theta_1 \leq \tilde{\theta}$

Without loss of generality, we assume $C_1 \geq C_2$. Let $\beta \triangleq \frac{C_1}{C_2}$, and hence we have $\beta \geq 1 - \beta$ (i.e., $1 \geq \frac{1-\beta}{\beta}$).

First, we consider partition-preferred congestion function $K(Q, C) \geq K(\alpha Q, \alpha C)$ for all $0 \leq \alpha \leq 1$. By Eqn. (34), it follows that

$$K(Q_2, C_2) = K(Q_1, C_1) \geq K\left(\frac{1-\beta}{\beta}Q_1, C_2\right) \quad (37)$$

Because $K(Q, C)$ is increasing in Q , we obtain

$$Q_2 \geq \frac{1-\beta}{\beta}Q_1 \Rightarrow Q_1 + Q_2 \geq \frac{Q_1}{\beta} \quad (38)$$

Because $K(Q, C) \geq K(\alpha Q, \alpha C)$ for all $0 \leq \alpha \leq 1$, we obtain

$$K(\tilde{Q}, C) \geq K(\beta \tilde{Q}, \beta C) = K(\beta \tilde{Q}, C_1) \quad (39)$$

We next use contradiction to show that $\theta_1 \geq \tilde{\theta}$. On the contrary, we suppose $\tilde{\theta} > \theta_1$. Then, by Eqns. (36) and (39), we obtain

$$K(Q_1, C_1) > K(\tilde{Q}, C) \geq K(\beta \tilde{Q}, C_1) \quad (40)$$

Hence, $Q_1 > \beta\tilde{Q}$ because $K(Q, C)$ is increasing in Q . Also, we note that

$$\tilde{\theta} > \theta_1 \Rightarrow F(\tilde{\theta}) > F(\theta_1) \Rightarrow \tilde{Q} > Q_1 + Q_2 \quad (41)$$

Therefore we derive

$$Q_1 > \beta\tilde{Q} > \beta(Q_1 + Q_2) \quad (42)$$

which is a contradiction to Eqn. (38). Hence it should be $\theta_1 \geq \tilde{\theta}$.

For multiplexing, to show that $K(Q, C) \leq K(\alpha Q, \alpha C)$ for all α implies $\theta_1 \leq \tilde{\theta}$, we note that we can reverse the signs “ \geq ” and “ $>$ ” in Eqns (37)-(42) to “ \leq ” and “ $<$ ” respectively.

Step 2: Next, we want to show:

1. if $\theta_1 \geq \tilde{\theta}$, then $S(p, p) \geq S(p)$ and $\pi(p, p) \geq \pi(p)$; and
2. if $\theta_1 \leq \tilde{\theta}$, then $S(p, p) \leq S(p)$ and $\pi(p, p) \leq \pi(p)$.

The case of provider profit $\pi(p, p)$ and $\pi(p)$ follows from Eqns. (10) and (13) and

$$\theta_1 \geq \tilde{\theta} \Rightarrow Q_1 + Q_2 \geq \tilde{Q} \Rightarrow \pi(p, p) \geq \pi(p) \quad (43)$$

$$\theta_1 \leq \tilde{\theta} \Rightarrow Q_1 + Q_2 \leq \tilde{Q} \Rightarrow \pi(p, p) \leq \pi(p) \quad (44)$$

where the service classes are charged at an identical price.

For the case of social welfare $S(p, p)$ and $S(p)$, by Eqn. (34), we obtain

$$S(p, p) = \int_0^{\theta_1} (V - \theta \cdot K(Q_2, C_2)) \cdot f(\theta) d\theta \quad (45)$$

From Eqn. (36), we obtain

$$\theta_1 \geq \tilde{\theta} \Rightarrow K(\tilde{Q}, C) \geq K(Q_2, C_2) \Rightarrow S(p, p) \geq S(p) \quad (46)$$

$$\theta_1 \leq \tilde{\theta} \Rightarrow K(\tilde{Q}, C) \leq K(Q_2, C_2) \Rightarrow S(p, p) \leq S(p) \quad (47)$$

which follows from Eqns. (11) and (44)-(47). \square

7.2 An Overview of Total Derivative

Before we proceed to the proof of Theorem 4.4, we briefly revisit the notion of total derivative [9] that will be useful in the following proofs.

The marginal change of real function $f(x, y, z)$ (i.e., derivative df) with respect to parameters x, y, z can be written as

$$df = \frac{\partial f}{\partial x} dx + \frac{\partial f}{\partial y} dy + \frac{\partial f}{\partial z} dz \quad (48)$$

where $\frac{\partial f}{\partial x}, \frac{\partial f}{\partial y}, \frac{\partial f}{\partial z}$ are the partial derivatives of f at the respective parameter while keeping other parameters as constants. Note that $\frac{\partial f}{\partial x}, \frac{\partial f}{\partial y}, \frac{\partial f}{\partial z}$ are also functions of x, y, z .

The derivative df can be regarded as a function of (dx, dy, dz) , each of them representing the marginal change of parameters x, y, z . Also, (dx, dy, dz) can be regarded as a vector in the 3D Euclidean space. An immediate consequence is that if function f is a stationary point at some (x_0, y_0, z_0) (e.g., the maximum), then $df = 0$ at (x_0, y_0, z_0) irrespective of what values of (dx, dy, dz) we pick. Otherwise, we will be possible to pick a vector (dx, dy, dz) , such that $df \neq 0$ implying that f is not a stationary point at (x_0, y_0, z_0) .

7.3 Proof for Theorem 4.4

Proof. First, we write $S = S(p_1, p_2)$. We study how the total derivative of S (see Sec. 7.2), dS , changes at $p_1 = p_2$. From Eqn. (12),

$$S = V \cdot F(\theta_1) - K(Q_2, C_2) \int_0^{\theta_2} \theta \cdot f(\theta) d\theta - K(Q_1, C_1) \int_{\theta_2}^{\theta_1} \theta \cdot f(\theta) d\theta \quad (49)$$

Recall that $k(Q_i, C_i) \triangleq \frac{\partial K(Q, C)}{\partial Q} \big|_{Q=Q_i, C=C_i}$, $Q_2 \triangleq F(\theta_2)$ and $Q_1 \triangleq F(\theta_1) - F(\theta_2)$.

Note that an equilibrium can be characterized by tuple (p_1, p_2) , a pair of independent variables. Similarly, an equilibrium can also be equivalently characterized by tuple (θ_1, θ_2) , by solving Eqn. (6) in Definition 3.1. Then (θ_1, θ_2) are treated as a pair of independent variables. Thus, we take the total derivative of S with respect to $(d\theta_1, d\theta_2)$, and obtain the following:

$$\begin{aligned} dS = & V \cdot f(\theta_1) d\theta_1 - \left(\int_0^{\theta_2} \theta \cdot f(\theta) d\theta \right) \cdot k(Q_2, C_2) f(\theta_2) d\theta_2 - K(Q_2, C_2) \cdot \theta_2 \cdot f(\theta_2) d\theta_2 \\ & - \left(\int_{\theta_2}^{\theta_1} \theta \cdot f(\theta) d\theta \right) \cdot k(Q_1, C_1) \left(f(\theta_1) d\theta_1 - f(\theta_2) d\theta_2 \right) - K(Q_1, C_1) \cdot \left(\theta_1 \cdot f(\theta_1) d\theta_1 - \theta_2 \cdot f(\theta_2) d\theta_2 \right) \end{aligned} \quad (50)$$

Then, at identical pricing $p_1 = p_2$, we have $K(Q_1, C_1) = K(Q_2, C_2)$. Hence, we obtain

$$\begin{aligned} dS|_{p_1=p_2} = & V \cdot f(\theta_1) d\theta_1 - \left(\int_0^{\theta_2} \theta \cdot f(\theta) d\theta \right) \cdot k(Q_2, C_2) f(\theta_2) d\theta_2 \\ & - \left(\int_{\theta_2}^{\theta_1} \theta \cdot f(\theta) d\theta \right) \cdot k(Q_1, C_1) \left(f(\theta_1) d\theta_1 - f(\theta_2) d\theta_2 \right) - K(Q_1, C_1) \cdot \theta_1 \cdot f(\theta_1) d\theta_1 \end{aligned} \quad (51)$$

Next, we pick a vector $(d\theta_1, d\theta_2)$, and show that $dS|_{p_1=p_2}$ will strictly increase in the direction of $(d\theta_1, d\theta_2)$. Such a vector indeed exists if we keep θ_1 as a constant (i.e., $d\theta_1 = 0$).

First, we obtain

$$\begin{aligned} dS|_{p_1=p_2, d\theta_1=0} & \quad (52a) \\ = & \left(\left(\int_{\theta_2}^{\theta_1} \theta \cdot f(\theta) d\theta \right) \cdot k(Q_1, C_1) - \left(\int_0^{\theta_2} \theta \cdot f(\theta) d\theta \right) \cdot k(Q_2, C_2) \right) f(\theta_2) d\theta_2 \\ > & \theta_2 \cdot \left(\left(\int_{\theta_2}^{\theta_1} f(\theta) d\theta \right) \cdot k(Q_1, C_1) - \left(\int_0^{\theta_2} f(\theta) d\theta \right) \cdot k(Q_2, C_2) \right) f(\theta_2) d\theta_2 \\ = & \theta_2 \left(Q_1 \cdot k(Q_1, C_1) - Q_2 \cdot k(Q_2, C_2) \right) f(\theta_2) d\theta_2 \end{aligned} \quad (52b)$$

Without loss of generality, we assume $Q_1 \geq Q_2$. Since the two service classes satisfy monotone preference (Definition 4.3), if we always pick $d\theta_2 > 0$ in the case of (m.1), and $d\theta_2 < 0$ in the case of (m.2), then it is always true that the total derivative $dS|_{p_1=p_2, d\theta_1=0} > 0$. Therefore, we see that the social welfare S can strictly increase by differentiated pricing ($p_1 \neq p_2$) from identical pricing ($p_1 = p_2$). \square

7.4 Proof for Corollary 4.5

Proof. By Theorem 4.4, it is true for $m = 2$. For $m = 3$, it follows that $p_1 = p_2 = p_3$ cannot be optimal. Next, we also show that $p_1 > p_2 = p_3$ and $p_1 = p_2 > p_3$ cannot be optimal. Then, the total

derivative of social welfare S with three service classes becomes

$$\begin{aligned}
dS = & \\
& V \cdot f(\theta_1)d\theta_1 - \left(\int_0^{\theta_3} \theta \cdot f(\theta)d\theta \right) \cdot k(Q_3, C_3)f(\theta_3)d\theta_3 - K(Q_3, C_3) \cdot \theta_3 \cdot f(\theta_3)d\theta_3 \\
& - \left(\int_{\theta_3}^{\theta_2} \theta \cdot f(\theta)d\theta \right) \cdot k(Q_2, C_2) \left(f(\theta_2)d\theta_2 - f(\theta_3)d\theta_3 \right) - K(Q_2, C_2) \cdot \left(\theta_2 \cdot f(\theta_2)d\theta_2 - \theta_3 \cdot f(\theta_3)d\theta_3 \right) \\
& - \left(\int_{\theta_2}^{\theta_1} \theta \cdot f(\theta)d\theta \right) \cdot k(Q_1, C_1) \left(f(\theta_1)d\theta_1 - f(\theta_2)d\theta_2 \right) - K(Q_1, C_1) \cdot \left(\theta_1 \cdot f(\theta_1)d\theta_1 - \theta_2 \cdot f(\theta_2)d\theta_2 \right)
\end{aligned} \tag{53}$$

For $p_1 > p_2 = p_3$ (i.e., $K(Q_2, C_2) = K(Q_3, C_3)$), setting $d\theta_1 = d\theta_2 = 0$ will degenerate to the case $m = 2$. For $p_1 = p_2 > p_3$ (i.e., $K(Q_1, C_1) = K(Q_2, C_2)$), setting $d\theta_1 = d\theta_3 = 0$ will also degenerate to the case $m = 2$. Hence, it follows for $m = 3$ is true. Using an iterative argument, we can show that it is true for all $m \geq 2$. \square

7.5 Proof for Theorem 4.6

Proof. Similar to Theorem 4.4, taking the total derivative of π with respect to (dp_1, dp_2) , we have

$$\begin{aligned}
d\pi = & Q_2 dp_2 + p_2 \cdot f(\theta_2) \left(\frac{\partial \theta_2}{\partial p_1} dp_1 + \frac{\partial \theta_2}{\partial p_2} dp_2 \right) + Q_1 dp_1 \\
& + p_1 \cdot \left(f(\theta_1) \left(\frac{\partial \theta_1}{\partial p_1} dp_1 + \frac{\partial \theta_1}{\partial p_2} dp_2 \right) - f(\theta_2) \left(\frac{\partial \theta_2}{\partial p_1} dp_1 + \frac{\partial \theta_2}{\partial p_2} dp_2 \right) \right)
\end{aligned} \tag{54}$$

Then, at identical pricing ($p_1 = p_2$), we obtain

$$d\pi|_{p_2=p_1} = Q_2 dp_2 + Q_1 dp_1 + p_1 \cdot f(\theta_1) \left(\frac{\partial \theta_1}{\partial p_1} dp_1 + \frac{\partial \theta_1}{\partial p_2} dp_2 \right) \tag{55}$$

However, $d\pi$ is more difficult than dS , involving $\frac{\partial \theta_1}{\partial p_1}$ and $\frac{\partial \theta_1}{\partial p_2}$.

As in Theorem 4.4, we pick a vector (dp_1, dp_2) , and show that $d\pi|_{p_1=p_2}$ will strictly increase in the direction of (dp_1, dp_2) . To achieve this we keep θ_1 as a constant (i.e., $d\theta_1 = \frac{\partial \theta_1}{\partial p_1} dp_1 + \frac{\partial \theta_1}{\partial p_2} dp_2 = 0$). Hence

$$d\pi|_{p_1=p_2, d\theta_1=0} = Q_1 dp_1 + Q_2 dp_2 \tag{56}$$

Also, from Eqn. (6) in Definition 3.1, we obtain the total derivatives (p_1, p_2) with respect to (dp_1, dp_2) as

$$\begin{cases} dp_1 = -K(Q_1, C_1)d\theta_1 - \theta_1 \cdot k(Q_1, C_1)(f(\theta_1)d\theta_1 - f(\theta_2)d\theta_2) \\ dp_1 - dp_2 = (K(Q_2, C_2) - K(Q_1, C_1))d\theta_2 \\ \quad + \theta_2 \cdot \left(k(Q_2, C_2)(f(\theta_2)d\theta_2) - k(Q_1, C_1)(f(\theta_1)d\theta_1 - f(\theta_2)d\theta_2) \right) \end{cases} \tag{57}$$

We keep θ_1 as a constant (i.e., $d\theta_1 = 0$), and by identical pricing $p_1 = p_2 \Rightarrow K(Q_1, C_1) = K(Q_2, C_2)$, we have

$$\begin{cases} dp_1 = \theta_1 \cdot k(Q_1, C_1)f(\theta_2)d\theta_2 \\ dp_1 - dp_2 = \theta_2 \cdot \left(k(Q_2, C_2) + k(Q_1, C_1) \right) \cdot f(\theta_2)d\theta_2 \end{cases} \tag{58}$$

Solving Eqn. (58) for (dp_1, dp_2) , we obtain

$$dp_2 = \frac{-\theta_2 \cdot (k(Q_2, C_2) + k(Q_1, C_1)) + \theta_1 \cdot k(Q_1, C_1)}{\theta_1 \cdot k(Q_1, C_1)} dp_1 \tag{59}$$

Because keeping θ_1 as a constant (i.e., $d\theta_1 = 0$), by substituting Eqn. (59) we obtain

$$d\pi|_{p_1=p_2, d\theta_1=0} = Q_2 dp_2 + Q_1 dp_1 \quad (60a)$$

$$= \left(Q_2 \cdot (-\theta_2 \cdot (k(Q_2, C_2) + k(Q_1, C_1)) + \theta_1 \cdot k(Q_1, C_1)) + Q_1 \theta_1 \cdot k(Q_1, C_1) \right) \frac{dp_1}{\theta_1 \cdot k(Q_1, C_1)}$$

$$= (-Q_2 \theta_2 \cdot (\frac{k(Q_2, C_2)}{k(Q_1, C_1)}) - Q_2 \theta_2 + Q_2 \theta_1 + Q_1 \theta_1) \frac{dp_1}{\theta_1} \quad (60b)$$

$$> (-Q_2 \cdot (\frac{k(Q_2, C_2)}{k(Q_1, C_1)}) - Q_2 + Q_2 + Q_1) \frac{\theta_2 dp_1}{\theta_1} \quad (60c)$$

$$= (Q_1 \cdot k(Q_1, C_1) - Q_2 \cdot k(Q_2, C_2)) \frac{\theta_2 dp_1}{\theta_1 \cdot k(Q_1, C_1)} \quad (60d)$$

Without loss of generality, we assume $Q_1 > Q_2$. Since the two service classes satisfy monotone preference (Definition 4.3), if we always pick $d\theta_2 > 0$ in the case of (m.1) and $d\theta_2 < 0$ in the case of (m.2), then it is always true that $d\pi|_{p_1=p_2, d\theta_1=0} > 0$. Therefore, the provider profit π can strictly increase by differentiated pricing ($p_1 \neq p_2$) from identical pricing ($p_1 = p_2$). \square

7.6 Derivatives of $\frac{d\pi^I}{dp^I}$

Table 2: Derivatives of all cases of $\frac{d\pi^I}{dp^I}$.

Cases	$\frac{d\pi^I}{dp^I}$
$p^I \geq p_1^I \geq p_2^I$	$\frac{1}{k_1 \theta_1} \left(-p^I + k_1 Q_1 \theta_1 \right.$
$p_1^I \geq p^I \geq p_2^I$	$\left. + \frac{K_1 p^I (-k_2 k_3 \theta_2 \theta_3 + k_1 \theta_1 (K_2 - K_3 - 2k_3 \theta_3) + (K_1 - K_2 - 2k_2 \theta_2)(K_2 - K_3 - 2k_3 \theta_3))}{k_1 k_2 \theta_1 \theta_2 (K_2 - K_3 - 2k_3 \theta_3) - (K_1 + k_1 \theta_1)(k_2 k_3 \theta_2 \theta_3 - (K_1 - K_2 - 2k_2 \theta_2)(K_2 - K_3 - 2k_3 \theta_3))} \right)$
$p_1^I \geq p_2^I \geq p^I$	$Q_2 + \frac{p^I}{K_2 - K_3 - 2k_3 \theta_3}$
	$- \frac{p^I (K_1 + k_1 \theta_1)(K_2 - K_3 + k_2 \theta_2 - 2k_3 \theta_3)(-K_2 + K_3 + k_3 \theta_3)}{(-K_2 + K_3 + 2k_3 \theta_3)(k_2 k_3 (K_1 + k_1 \theta_1) \theta_2 \theta_3 + (-k_1 k_2 \theta_1 \theta_2 - (K_1 + k_1 \theta_1)(K_1 - K_2 - 2k_2 \theta_2))(K_2 - K_3 - 2k_3 \theta_3))}$
	$Q_3 + \frac{p^I}{K_2 - K_3 - 2k_3 \theta_3}$
	$+ \frac{k_2 k_3 p^I (K_1 + k_1 \theta_1) \theta_2 \theta_3}{(-K_2 + K_3 + 2k_3 \theta_3)(k_2 k_3 (K_1 + k_1 \theta_1) \theta_2 \theta_3 + (-k_1 k_2 \theta_1 \theta_2 - (K_1 + k_1 \theta_1)(K_1 - K_2 - 2k_2 \theta_2))(K_2 - K_3 - 2k_3 \theta_3))}$
$p^I \geq p^I$	$\frac{K_1^2 Q_1 + K_2(p^I - Q_1(K_1 + k_1 \theta_1)) + k_2(p^I - k_1 Q_1 \theta_1) \theta_2 + K_1 Q_1 (k_1 \theta_1 - 2k_2 \theta_2)}{K_1^2 - k_1 \theta_1 (K_2 + k_2 \theta_2) - K_1 (K_2 - k_1 \theta_1 + 2k_2 \theta_2)}$
$p^I \geq p^I$	$Q_2 - \frac{p^I (K_1 + k_1 \theta_1)}{-k_1 k_2 \theta_1 \theta_2 - (K_1 + k_1 \theta_1)(K_1 - K_2 - 2k_2 \theta_2)}$